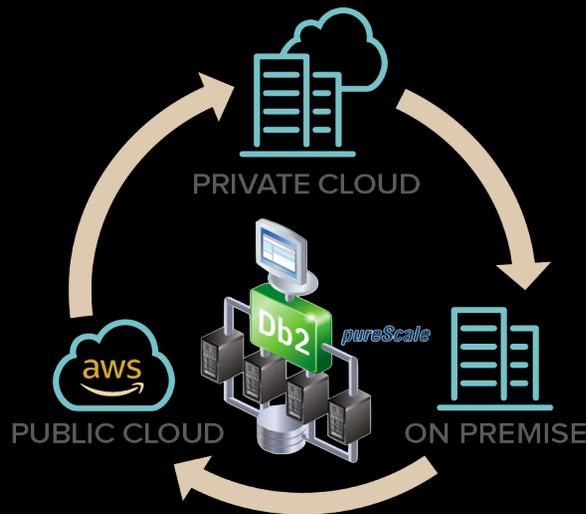


Db2 pureScale on AWS

The New Experience

The Db2 Night Show™



- pureScale refresher
- The New Offering
- The Architecture
- Support
- Roadmap

Alan Lee

Program Director, Data & AI, Db2 LUW Development



ykalee@ca.ibm.com



@ykalee1



www.linkedin.com/in/ykalee

Please note

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

NOTICE AND DISCLAIMER

- © 2022 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity. IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

pureScale refresher

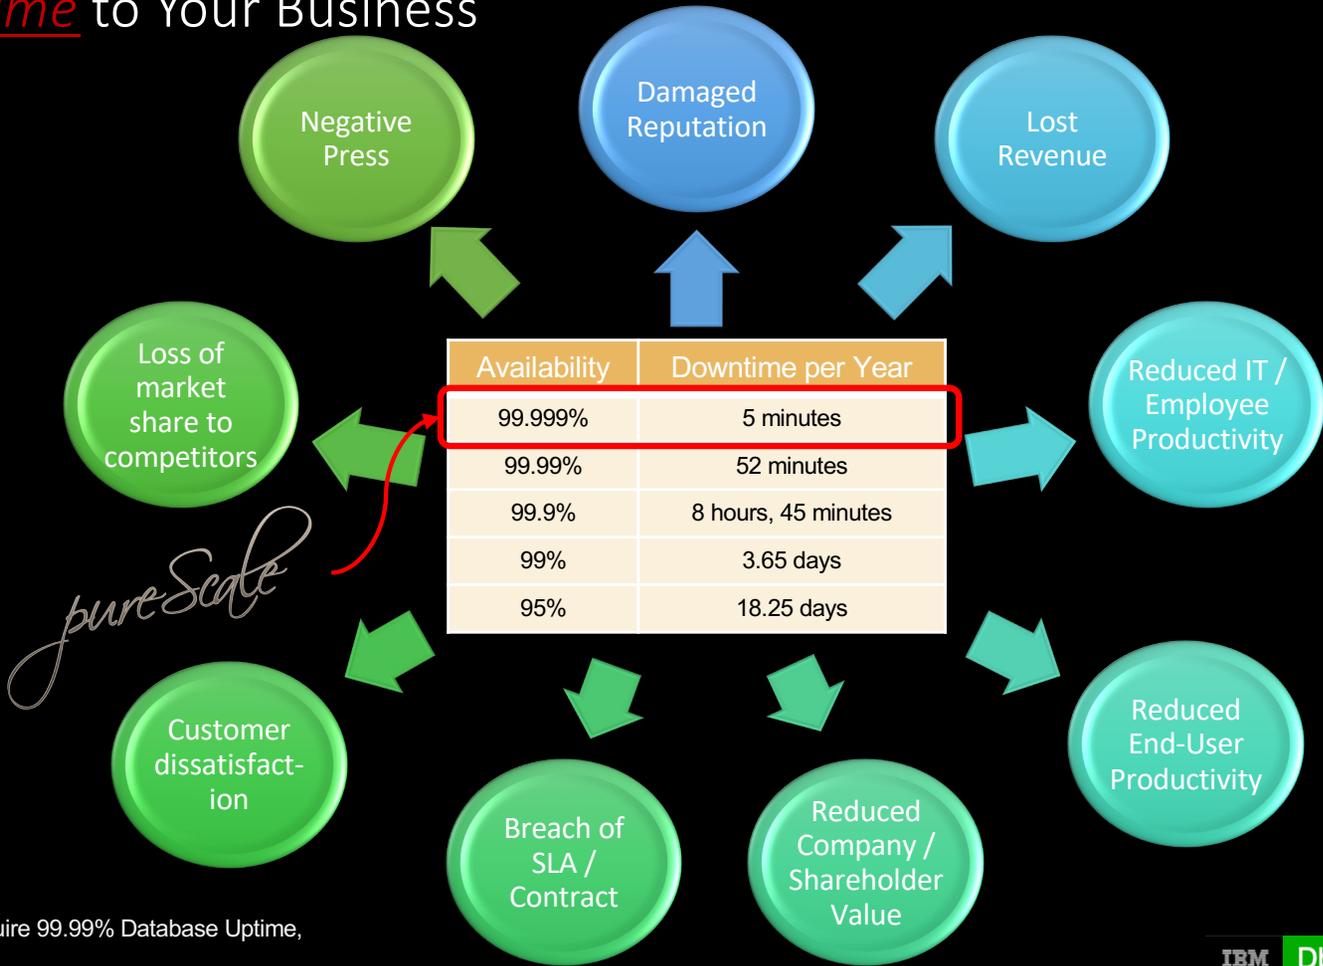
- The “Why”
- The “What”
- The deployment base



The Impact of Downtime to Your Business

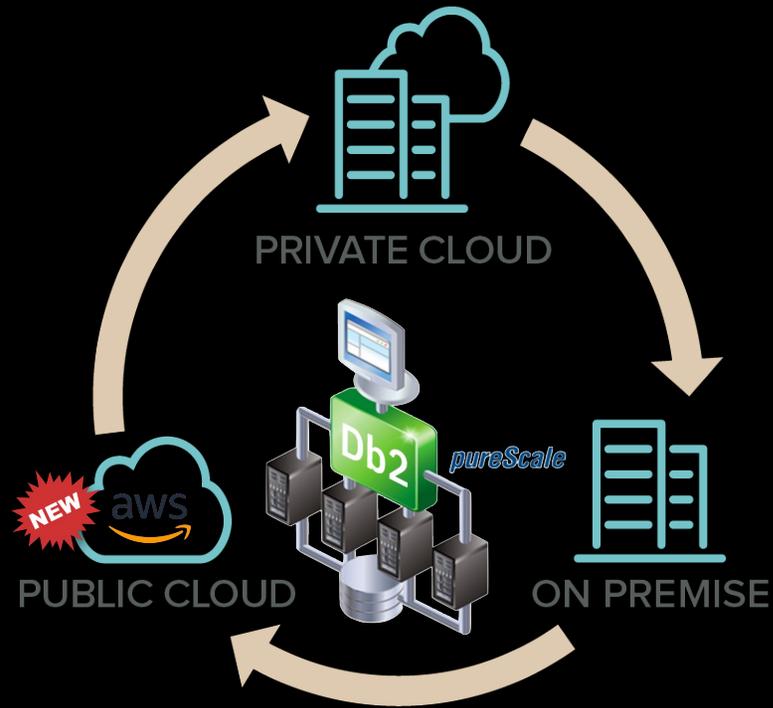


64% of organizations require that their databases deliver a minimum of 99.99% or better uptime for their most mission critical applications *



* ITIC Paper "Two-Thirds of Corporations Now Require 99.99% Database Uptime, Reliability ", Laura DiDio, July 10th, 2013

Db2's five 9s solution is pureScale

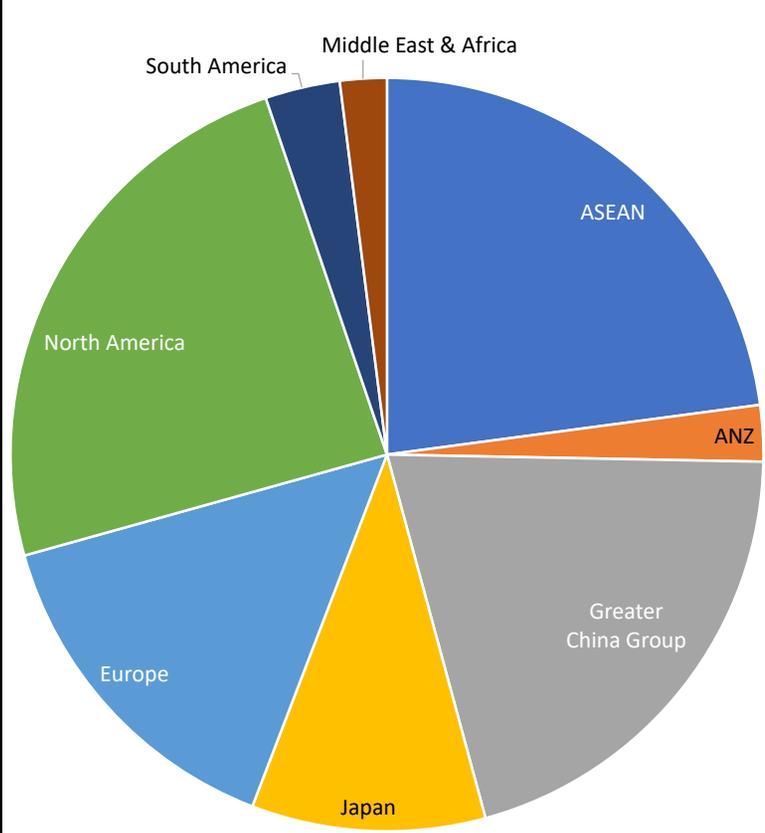


- Elastic Capacity
- Extreme Availability
- Application Transparency
- Workload Consolidation
- NEW** Cloud-Ready

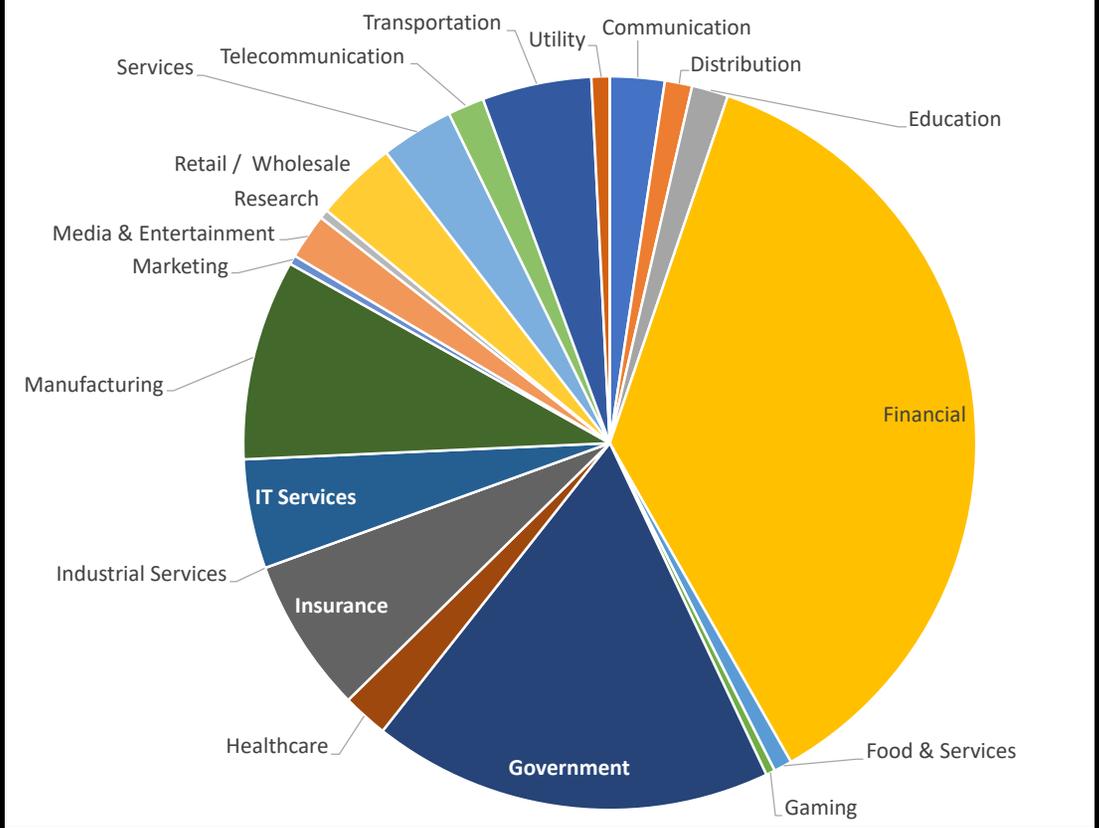
Learning from the undisputed Gold Standard... System z

pureScale Successes at Every Geography and Every Industry

Db2 pureScale Customers - World Wide



Db2 pureScale Industries - World Wide



The New Offering

- AWS Marketplace landing page
- Our documentation
- The Deployment Journey
- Value Propositions
- T-shirt size and guidance



The Deployment Guide in IBM Documentation

AWS Marketplace, leveraging the *Bring Your Own License* model. For more information, download the guide [Deploying IBM® Db2 pureScale with TCP/IP private network on AWS](#).' The left sidebar shows a navigation menu with 'Db2 pureScale on AWS' selected. The top navigation bar includes the IBM logo, 'Documentation', and a search bar with 'Search in Db2 11.5'."/>

IBM Documentation Search in Db2 11.5

Db2 / 11.5 /

Db2 pureScale on Amazon Web Services (AWS)

Last Updated: 2022-07-22

You can run a Db2 pureScale cluster as a self-managed service on AWS.

This self-managed service is available through [AWS Marketplace](#), leveraging the *Bring Your Own License* model.

For more information, download the guide [Deploying IBM® Db2 pureScale with TCP/IP private network on AWS](#).

Parent topic:
→ [Db2 for cloud service providers](#)

Goal:

- move all content in the PDF directly into IBM Documentation in future release

Table of Contents	
Introduction.....	3
Value Propositions with pureScale on AWS	4
Extreme capacity.....	4
Continuous availability.....	4
Application transparency.....	4
Use cases.....	5
Cluster readiness from weeks/months to minutes.....	5
Simplified Redeployment.....	5
"Fastpass" to new Db2 levels validation with dramatic lower TCO.....	6
Lower cost with S3.....	6
Disaster Recovery with extreme ease.....	7
Transferable skills.....	7
Deployment Overview.....	7
Single Availability Zone (AZ) Deployment Model.....	7
Multiple Availability Zones (AZs) Deployment Model.....	9
Supported Regions.....	9
Prerequisites and Requirements.....	10
Technical prerequisites for deployment.....	10
Host.....	10
Operating System Level.....	10
Db2 version.....	10
TSA, RSCT, and Spectrum Scale levels.....	10
Cluster interconnect.....	10
Network Interfaces.....	10
Storage.....	10
Skills and specialized knowledge.....	10
Environment configuration requirements.....	11
AWS account.....	11
Resource quotas.....	11
Architecture Diagrams.....	12
Security.....	13
Root Privilege.....	13
IAM roles for administrating and using the Db2 pureScale deployment.....	13
S3 buckets access.....	13
Authentication keys used by the deployment process.....	13
Costs.....	14
AWS services.....	14
Software licenses.....	14
Sizing.....	14
Performance.....	15
Deployment Assets.....	15
Deployment Steps.....	15
Set Up Client Connectivity via AWS Transit Gateway.....	19
Health Check.....	20
Backup and Recovery.....	20
Routine Maintenance.....	20
Emergency Maintenance.....	20
Support.....	21
FAQ.....	21

The offering in AWS Marketplace

IBM Db2 pureScale
By: [IBM Data and AI](#) Latest Version: 1

The IBM® Db2® pureScale® on AWS feature eliminates unplanned downtime by providing your mission critical workloads with continuous availability, application transparency, and shared

Linux/Unix **BYOL**

Continue to Subscribe

Save to List

Typical Total Price
\$2.165/hr
Total pricing per instance for services hosted on c6i.12xlarge in US East (N. Virginia). [View Details](#)

Overview Pricing Usage Support Reviews

Product Overview

Mission critical workloads require continuous availability at scale. With best-in-class failure detection and recovery, eliminate unplanned downtime while achieving high performance with IBM® Db2® pureScale®. pureScale® leverages the IBM® Db2® parallel sysplex architecture, providing mainframe-class continuous availability that runs anywhere, whenever you need it. pureScale® enables shared cluster scale out, starting with a small cluster, and allowing you to grow your workloads as much or as little as you need. With IBM® Db2® pureScale® there are no application changes to make, data to redistribute, or performance tuning to do. Provide application transparency through automatic load balancing across all active members. Achieve optimal resource utilization at all times, which helps to keep application response times low, while reducing the risk and cost of application changes. Start building more, while managing less.

Highlights

- Cluster of any size ready in minutes.
- Dramatic reduction in total cost of ownership.
- Disaster Recovery with extreme ease.

Version	1
By	IBM Data and AI
Categories	Databases & Analytics Platforms Databases
Operating System	Linux/Unix, SUSE 12 SP5
Delivery Methods	CloudFormation Template

Pricing Information

Use this tool to estimate the software and infrastructure pricing for services hosted in your region. Your usage and costs might be different from this estimate. They will be reflected on your monthly AWS billing reports.

Estimating your costs

Choose your region and fulfillment option to see the modified estimated price by choosing different instances.

Region: **US East (N. Virginia)**

Fulfillment Option: Db2 pureScale on AWS

Software Pricing Details

Db2 pureScale **\$0 /hr** running on c6i.12xlarge

Infrastructure Pricing Details

Estimated Infrastructure Cost: **\$11205/month using 5x c6i.12xlarge instance(s)**

BYOL Available for customers with current licenses purchased via other channels.

The table shows current software and infrastructure pricing for services hosted in US East (N. Virginia). Additional taxes or fees may apply. Use of Local Zones or Wavelength infrastructure deployment may alter your final pricing.

EC2 Instance type	Software/hr	EC2/hr	Total/hr
<input type="radio"/> c6i.4xlarge	\$0	\$0.805	\$0.805
<input type="radio"/> c6i.8xlarge	\$0	\$1.485	\$1.485
<input checked="" type="radio"/> c6i.12xlarge	\$0	\$2.165	\$2.165
<input type="radio"/> c6i.16xlarge	\$0	\$2.845	\$2.845

Usage Information

Fulfillment Options

Db2 pureScale on AWS
CloudFormation Template

This pureScale cluster deployment sets up three Db2 members (five for the largest deployments) and two CFs, in a newly created VPC. Applications can connect to the pureScale cluster from within this VPC, from other VPCs or via VPN through the VPC transit gateway.

- View Template Components
- View usage instructions
- View CloudFormation Template

End-user license agreement

By subscribing to this product you agree to terms and conditions outlined in the product [End User License Agreement \(EULA\)](#)

Additional Resources

- [Introduction to Db2 pureScale](#)
- [Overview of Db2 pureScale on AWS](#)
- [FAQ for Db2 pureScale on AWS \(see end of document\)](#)

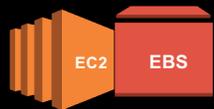
CloudFormation Template

AWS CloudFormation templates are JSON or YAML formatted text files that simplify provisioning and management on AWS. The templates describe the service or application architecture you want to deploy and AWS CloudFormation uses those templates to provision and configure the required services (such as Amazon EC2 instances or Amazon RDS DB instances). The deployed application and associated resources is called a "stack". [Learn more](#)

User Experience (without Automation)

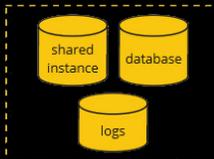
1

Provisioning AWS Resources



2

Configure Multi-Attach Volumes



Deployment Guide – Configuring AWS Resources

- 1 x Virtual Private Cloud (VPC)
 - IPv4 CIDR 10.0.0.0/16
- 1 x Internet Gateway
- 1 x Route Table
- 2 x Subnets
 - One public subnet with IPv4 CIDR 10.0.0.0/24
 - One private subnet with IPv4 CIDR 10.0.1.0/24
- 1 x Network Access Control List (ACL)
- 1 x Security Group
 - * Allow traffic between public+private subnets
 - * SSH port open
- 5 x EC2 Instances, with the following per instance:
 - 1 x EBS Volume
 - 1 x Elastic IP
 - 2 x Network Interfaces
 - One part of public subnet
 - One part of private subnet
 - EBS-Optimized
 - Clustered Placement Group
- 4 x EBS Multi-Attach Volume
 - (Optional) 4 x EBS Snapshot

Deployment Guide – Installing Db2

1. Configure root access (ssh) for EC2 instances
2. Generate ssh key pairs for passwordless ssh
3. (Optional) Change root password for instances
4. Install and enable multipath
5. Configuring /etc/hosts
6. Enabling passwordless ssh
7. Update zypper repositories
8. Installing required packages for Db2: `patch make libgomp1 kernel-source kernel-headers m4 gcc-c++ cpp gcc chrony mksh`
9. Rebooting All Hosts
10. Configuring Known Hosts
11. Transferring Install Image to all hosts (BYOL)
12. Extracting Install Image on all hosts
13. Installing Db2 on all hosts (Db2_install)
14. Creating Spectrum Scale Cluster (mmcrcluster)
15. Accepting Spectrum Scale License
16. Creating Spectrum Scale NSD (Stanza, mmcrnsd)
17. Starting Spectrum Scale (mmstartup)
18. Creating Spectrum Scale File System (mmcrfs)
19. Creating Db2 Users and Groups
20. Creating Db2 Instances (db2icrt, db2iupdt)

3

Installing Db2 pureScale



4

Setup database/ Restore existing



Total Time - Approximately 2-3 Hours

User Experience (with Automation)

1

Specify Non-Default Deployment Parameters

1. Deployment Type: **XS/S/M/L**
2. Root password
3. Availability Zone
4. SecurityGroupCIDR: **0.0.0.0/0**
5. SSH Keypair: **mykey.pem**



~ 36 AWS Resources + Configuration

Seamless Db2 pureScale + Spectrum Scale Installation

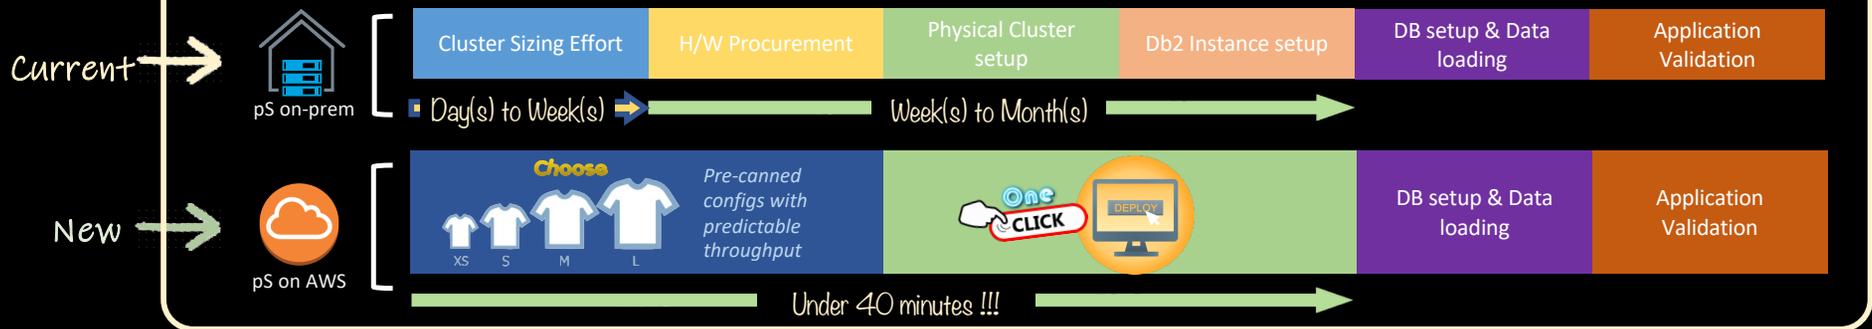
2

Setup database/ Restore existing



Total Time - Approximately 40 Minutes for 3 members and 2 CFs

Up & Running *In Minutes*



A Revolutionary Experience



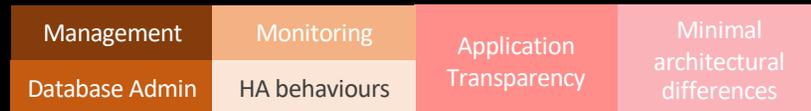
Much Lower TCO



with new and/or existing use cases



Transferrable skills from on-premises experiences



Deployment Cluster size and throughput guidance

X-SMALL (c6i.4xlarge)	SMALL (c6i.8xlarge)	MEDIUM (c6i.12xlarge)	LARGE (c6i.16xlarge)
3 Member + 2 CFs	3 Member + 2 CFs	3 Member + 2 CFs	5 Member + 2 CFs
❖ 16 CPU ❖ 32 Gb RAM ❖ 5000 IOPS ❖ Up to 12.5 GbE	❖ 32 CPU ❖ 64 Gb RAM ❖ 5000 IOPS ❖ 12.5 GbE	❖ 48 CPU ❖ 96 Gb RAM ❖ 7000 IOPS ❖ 18.75 GbE	❖ 64 CPU ❖ 128Gb RAM ❖ 10000 IOPS ❖ 25 GbE

Application Read/Write Ratio	SQL Statements Per Second			
70/30	70,000	125,000	300,000	350,000
90/10	110,000	190,000	440,000	770,000

**On-premises
4 Member + 2 CFs**

- ❖ 20 CPU
- ❖ 192 Gb RAM
- ❖ Comparable IOPS
- ❖ 10 GbE

R/W = 70/30 - yields 100,000 SQL statements per second

Instance setting, conventions

Example of a XS, S, and M deployment:

```
db2inst1@ykalee-oregon-1:~> db2instance -list
```

ID	TYPE	STATE	HOME_HOST	CURRENT_HOST	ALERT	PARTITION_NUMBER	LOGICAL_PORT	NETNAME
0	MEMBER	STOPPED	ykalee-oregon-3	ykalee-oregon-3	NO	0	0	ykalee-oregon-3-eth1
1	MEMBER	STOPPED	ykalee-oregon-4	ykalee-oregon-4	NO	0	0	ykalee-oregon-4-eth1
2	MEMBER	STOPPED	ykalee-oregon-5	ykalee-oregon-5	NO	0	0	ykalee-oregon-5-eth1
128	CF	STOPPED	ykalee-oregon-1	ykalee-oregon-1	NO	-	0	ykalee-oregon-1-eth1
129	CF	STOPPED	ykalee-oregon-2	ykalee-oregon-2	NO	-	0	ykalee-oregon-2-eth1


```
db2inst1@ykalee-oregon-1:~>
```

HOSTNAME	STATE	INSTANCE_STOPPED	ALERT
ykalee-oregon-5	ACTIVE	NO	NO
ykalee-oregon-4	ACTIVE	NO	NO
ykalee-oregon-2	ACTIVE	NO	NO
ykalee-oregon-1	ACTIVE	NO	NO
ykalee-oregon-3	ACTIVE	NO	NO

netnames

- Append “-eth<#>” to the hostname.

hostnames

- User specifies a common base name of all hostnames as one of the give mandatory parameters before deployment
- Deployment appends “-<#>” to it starting with 1 for CF 128, followed by 2 for CF 129. First member starts with 3 and so on.

/etc/hosts

```
#Db2 pureScale Hostnames for Public Subnet
10.0.0.10 ykalee-oregon-1.us-west-2.compute.internal ykalee-oregon-1
10.0.0.11 ykalee-oregon-2.us-west-2.compute.internal ykalee-oregon-2
10.0.0.12 ykalee-oregon-3.us-west-2.compute.internal ykalee-oregon-3
10.0.0.13 ykalee-oregon-4.us-west-2.compute.internal ykalee-oregon-4
10.0.0.14 ykalee-oregon-5.us-west-2.compute.internal ykalee-oregon-5

#Db2 pureScale Hostnames for Private Subnet
10.0.1.10 ykalee-oregon-1-eth1.us-west-2.compute.internal ykalee-oregon-1-eth1
10.0.1.11 ykalee-oregon-2-eth1.us-west-2.compute.internal ykalee-oregon-2-eth1
10.0.1.12 ykalee-oregon-3-eth1.us-west-2.compute.internal ykalee-oregon-3-eth1
10.0.1.13 ykalee-oregon-4-eth1.us-west-2.compute.internal ykalee-oregon-4-eth1
10.0.1.14 ykalee-oregon-5-eth1.us-west-2.compute.internal ykalee-oregon-5-eth1
```

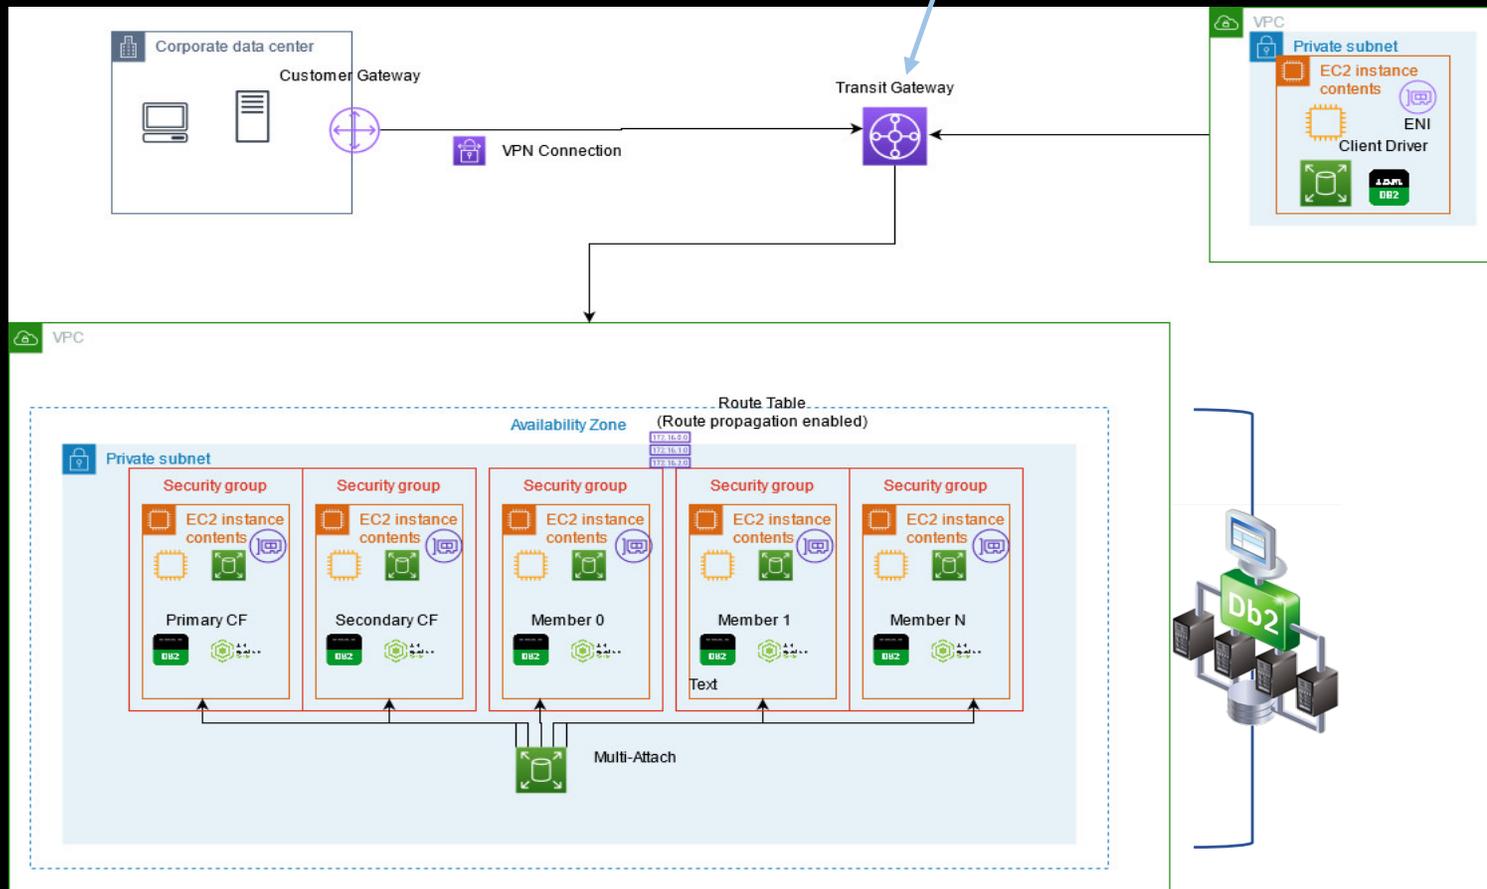
The Architecture

- 100,000' view
- On-premises Vs Cloud
- Disaster Recovery solution



Architecture Diagram (1 AZ)

AWS entity, setup by client following deployment guide



Same client setup as on-premises db2dsdrivers.cfg, etc.

One-Click Deployment

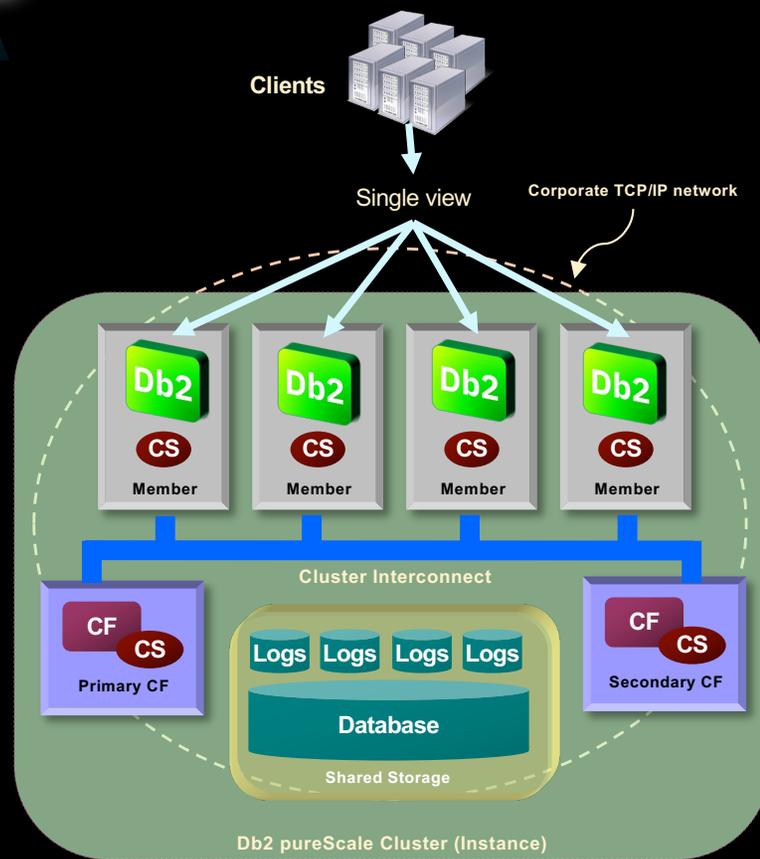
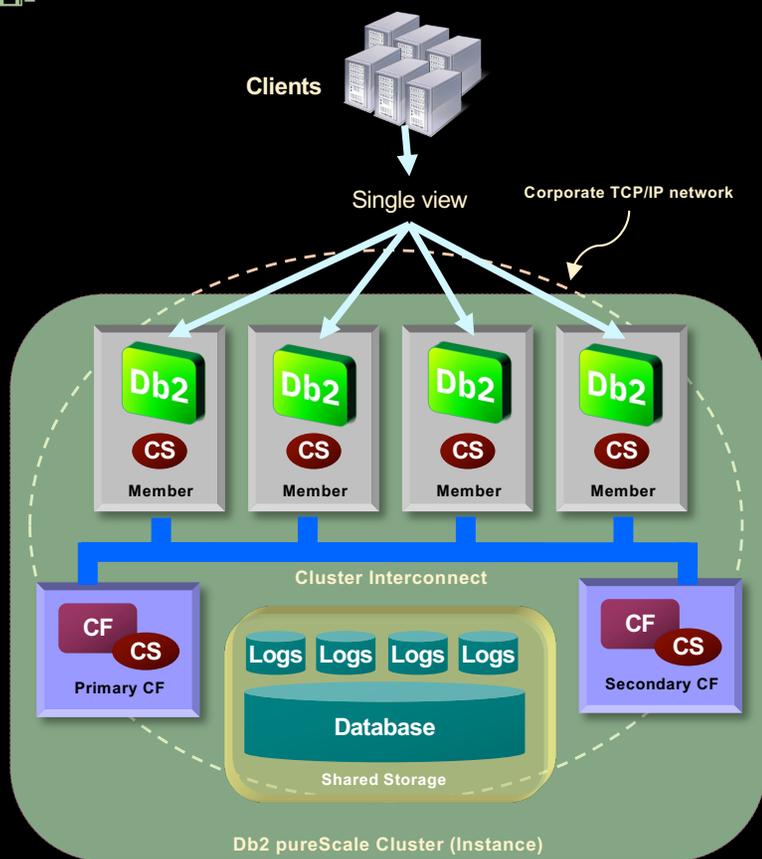
On-premise



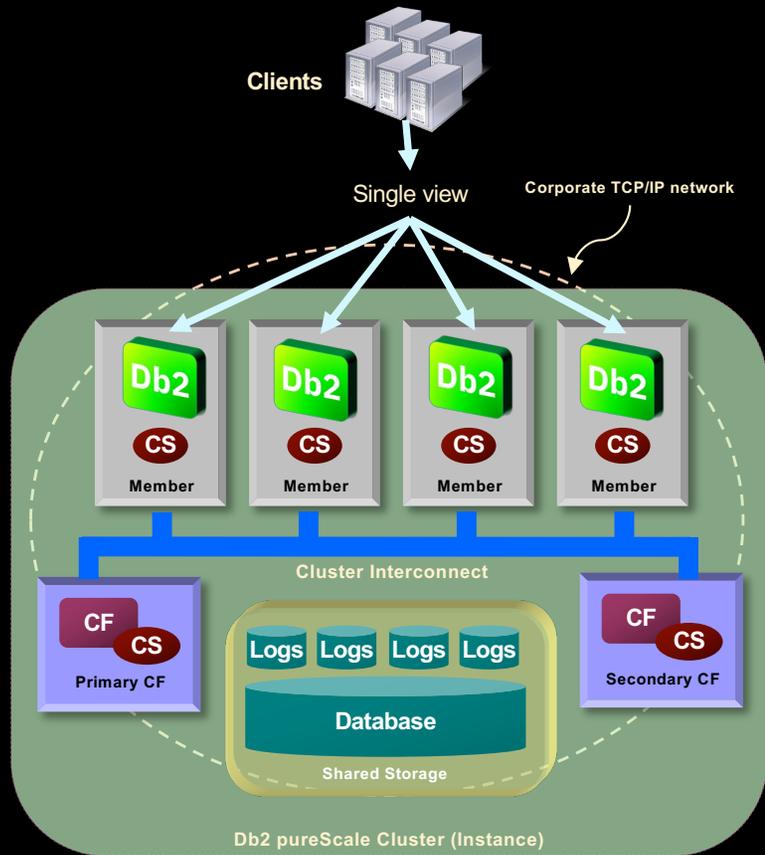
On-premises Architecture



Cloud Architecture



Component Side-by-Side View



* On selected architecture(s) only

On-premise	AWS
Client connect anywhere ... see single database	
<ul style="list-style-type: none"> • Clients connect into any member • Automatic workload balancing 	<ul style="list-style-type: none"> • No change functionally and Db2 setup • Need to setup AWS Transit gateway for inbound traffic outside of current VPC
Db2 engine runs on several host computers	
<ul style="list-style-type: none"> • Co-operate with each other to provide coherent access to the database from any member 	
Integrated Cluster Services (CS)	
<ul style="list-style-type: none"> • Failure detection, recovery automation, cluster FS • In partnership with STG (Spectrum Scale, RSCT) and Tivoli SA MP (TSA) (<i>to be replaced by Pacemaker later</i>) 	
Cluster Interconnect between members and CFs	
<ul style="list-style-type: none"> • Special optimizations provide low latency, high speed on RDMA-capable interconnects (e.g. 10/40*/100Gb RoCE) • Vanilla Ethernet also supported 	<ul style="list-style-type: none"> • Support Ethernet <i>for now</i> • RDMA equivalent in future release
Cluster caching facility (CF)	
<ul style="list-style-type: none"> • Efficient global locking and buffer management • Synchronous duplexing to secondary ensures availability 	
Data sharing architecture	
<ul style="list-style-type: none"> • Shared access to database • Members write to their own logs • Logs accessible from another host (used during recovery) 	<ul style="list-style-type: none"> • Same shared storage via new multi-attach EBS support in AWS • Limited to 16 host and within same region • No support of SCSI-3 PR.

Key “Delta” between offering on cloud vs on-premises



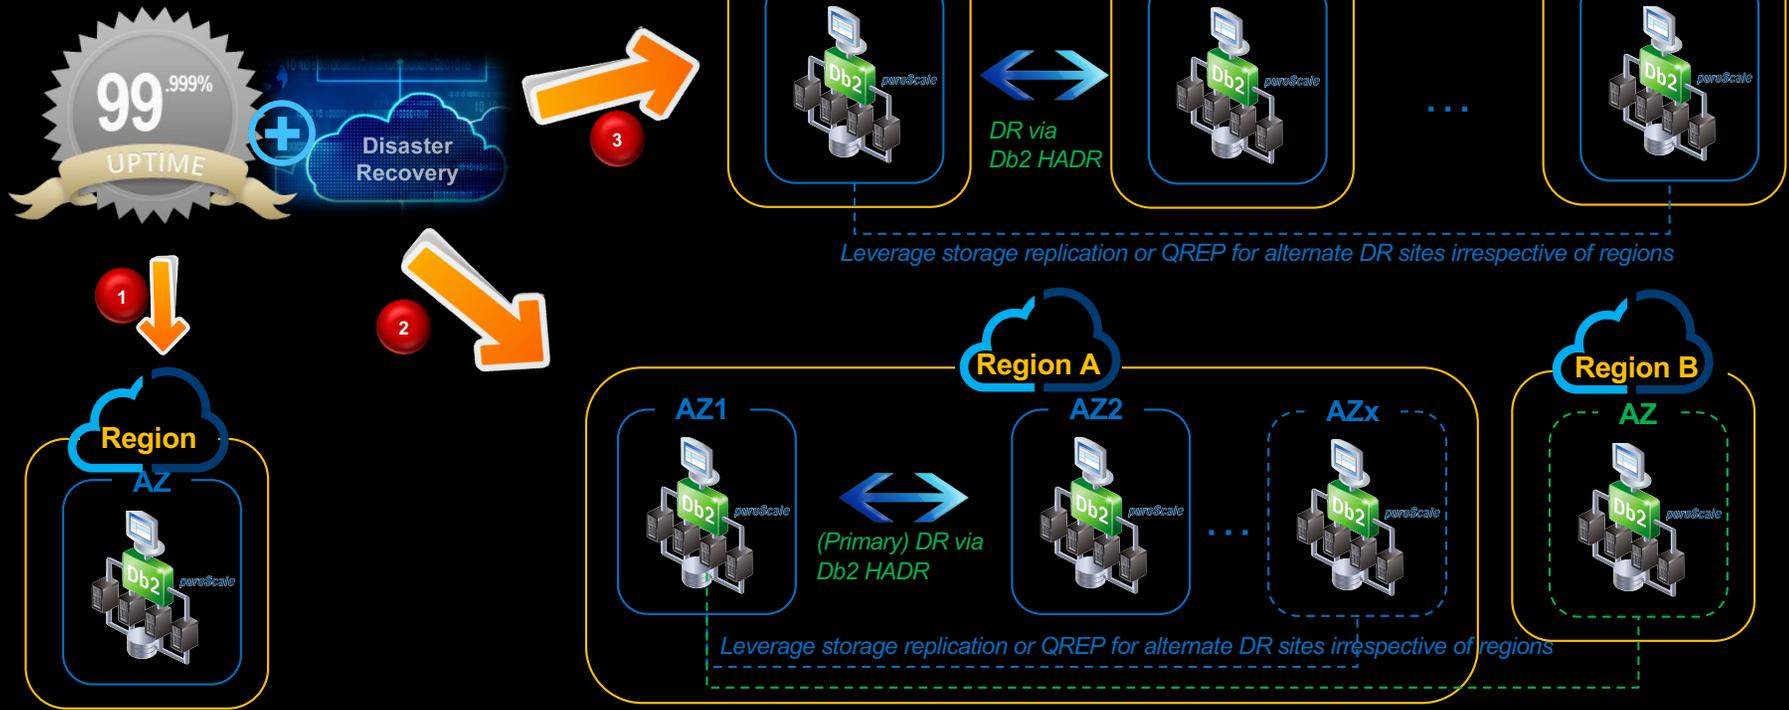
Speed, ease of deployment	Strength	Full autonomy, customized config.
No up-front cost	Capital Investment	Huge upfront + depreciation cost

Latest within same major release	Db2 version	“Any”
Intel	H/W Architecture	POWER, Intel, Z
RHEL or SLES	OS Choices	AIX, RHEL or SLES

Pre-configured size	Cluster Topology	Customizable
TCP/IP	Cluster Interconnect	RDMA, TCP/IP
Majority Node	Quorum	Disk tiebreaker
Disk Leasing	I/O Fencing	SCSI-3 PR
HADR	Integrated DR options	HADR, GDPC

The five 9s with DR cloud solution

GDPC



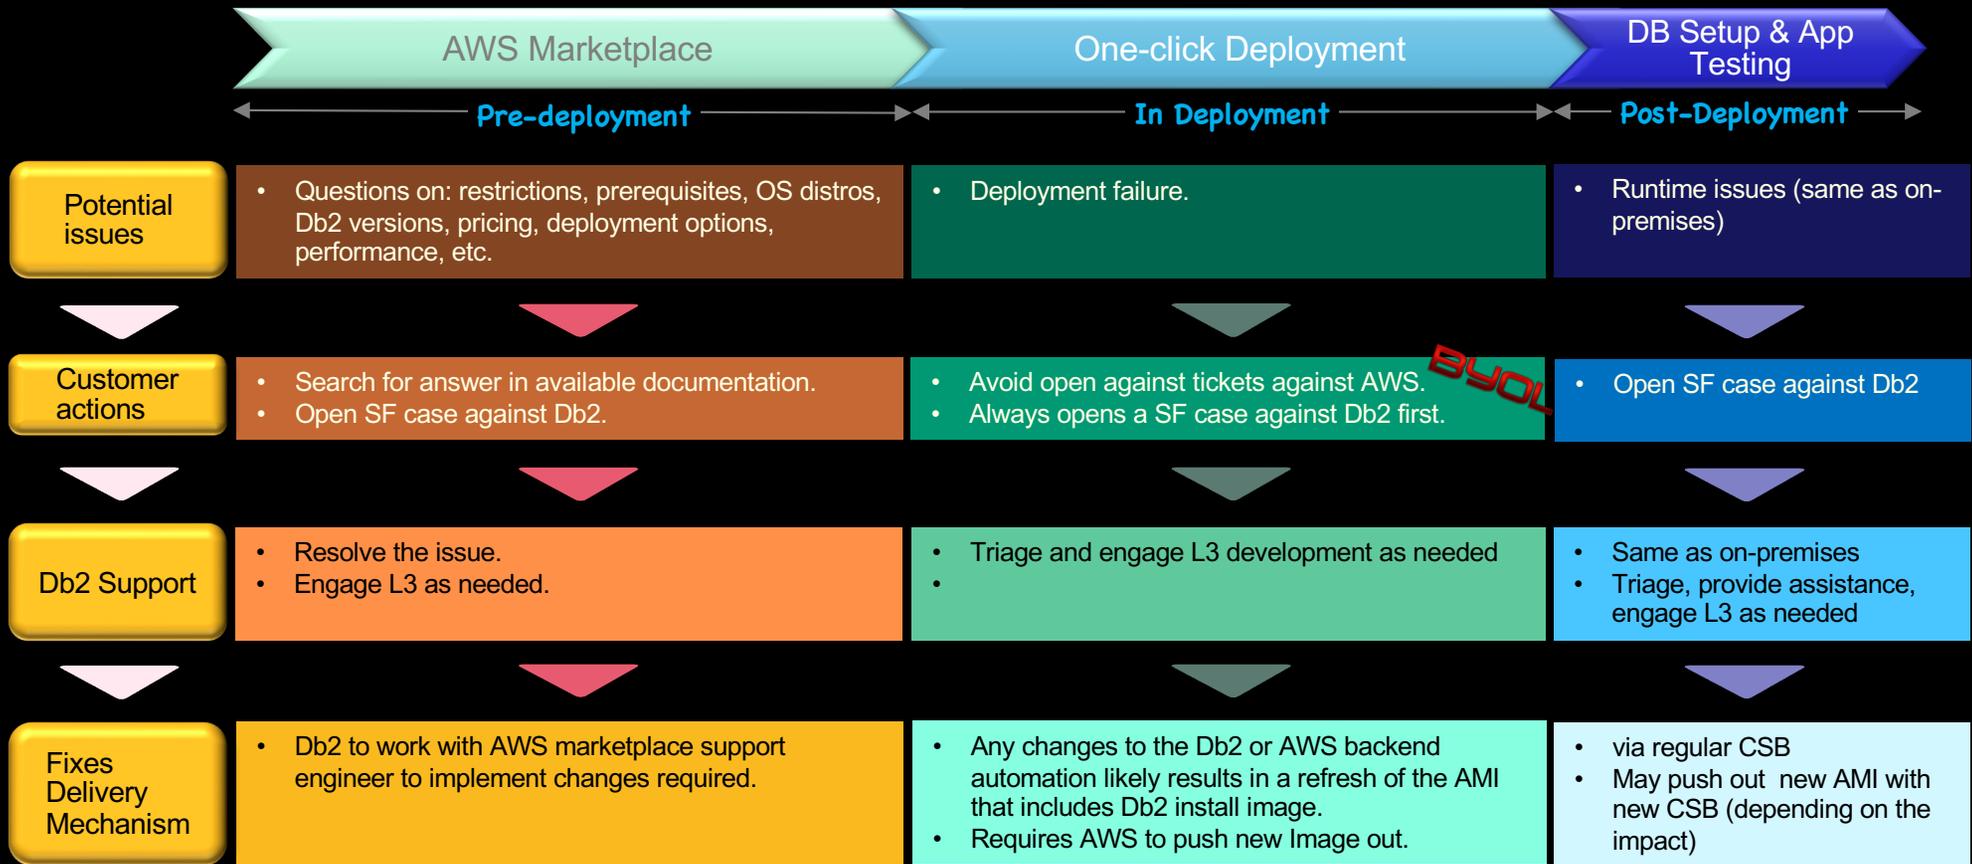
Goal: Leverage on-premises architecture as much as possible, preserve HA characteristics.

Support

- The 3 phases
- Common Issues
- Ongoing maintenance / updates



The three phases



Common Questions during Pre-deployment



Q1: The current offering is based on TSA as cluster manager, can we deploy Pacemaker in such a way to automate the failover from primary site to standby site ?

A: No

Q2: Can AWS EBS Snapshot be used in DR site (in a different region) ?

A: Yes, will need to suspend I/O just like on-premises via SET WRITE SUSPEND, etc.

Q3: Can altering the configuration of the EC2 instance after deployment, such as lowering IOPS at night for cost reduction, void support from IBM?

A: No, the cluster is under customers' control once deployed just as on-premises. But customers should consider the impact of change, validate it on a test environment before implementing it in production.

Q4: How many versions of Db2 will be kept in the AWS marketplace and how often will it be refreshed ?

A: Only the latest mod pack / fix pack in each major release (starting from 11.5) will be kept. It will always be refreshed at each fix pack, mod pack, and major release boundary. It may also be refreshed due to PSIRTS or certain pervasive HIPER APAR(s).

Note: These will be added to FAQ in [deployment guide](#)

Post-Deployment maintenance

Host maintenance for H/W, non-Db2 S/W, or OS upgrade:

- Follow “[Maintenance in a Db2 pureScale environment](#)”



Db2 fix pack, mod pack or cumulative special build update:

- Follow “[Applying fix packs in Db2 pureScale environments](#)” for:
 - Offline update procedures
 - Online update procedures
 - Concurrent update procedures

Changes to host configuration in terms of memory or cores may be dynamic from host perspective, but

Reminder!

- Db2 may not recognize it until after the next instance restart
- Db2 may not recognize it at all IF the memory size is **not** set to automatic

- **Placing mount points associated with database-level mount resources into maintenance mode**
You can perform maintenance on a specific mount point that is associated with a database level mount resource in the TSA resource model. Maintenance can be done without affecting the member resource's availability. This command must be run as the Db2 instance owner.
- **Performing maintenance on a Db2 pureScale host**
You can perform maintenance on or apply updates to a member host without affecting the availability of the databases in the Db2 pureScale instance.
- **Replacing both cluster caching facilities**
You can replace both cluster caching facilities using a rolling upgrade technique which ensures that your Db2 pureScale instance will not experience an outage.
- **Adding a disk to the shared file system**
After creating your shared file systems, there might be circumstances which require adding additional disks to the file system.
- **Quiescing a member**
Certain circumstances might require you to temporarily remove a member from the cluster (for example, maintenance operations).
- **Putting hosts into maintenance mode**
The term *host* referenced in this page refers to a Db2 member or Db2 cluster facility (CF) residing in either a physical host or a logical partition (LPAR) on supported platforms such as AIX. If you are applying software updates to Db2 cluster services or making changes to the server that impacts Db2, you must put the target host into maintenance mode. You can also put a host into maintenance mode if you want to ensure that members or cluster caching facilities are not restarted on the host when you are making updates to the operating system or hardware on the host.
- **Putting a cluster into maintenance mode**
You can put a cluster into maintenance mode when you are making updates to the operating system or hardware on the hosts in the cluster.
- **Moving a Db2 member or a cluster caching facility**
There are a number of reasons to move a Db2 member or a cluster caching facility from one host to another. This task outlines a number of possible scenarios and some additional factors to consider.
- **Recovering from a failed db2iupdt -add or -drop operation in a Db2 pureScale environment**
There are a number of reasons why a **db2iupdt -add** or **-drop** operation might fail, leaving the Db2 pureScale environment topology in an inconsistent state.
- **Clearing resource model alerts**
In a Db2 pureScale environment, hardware- or software-related alerts can exist on the status table for members, cluster caching facilities (CFs), and for hosts. In some cases, the alert conditions are transient so the alert field might clear itself; however, in other cases, the alert field remains set until the administrator resolves the problem and manually clears the alert.

Roadmap



Roadmap from feature perspective

Time

Categories	GA #1 (Q2 2022)	GA #2 (Q4 2022)	Future
Up & Running	One-click deployment from H/W to S/W		
Costs	AWS resources + Db2 license (BYOL)		
Disaster Recovery	HADR within same region; Qrep, storage	HADR - cross regions	
Db2 Version	V11.5.6.0	V11.5.8.0	Latest
Distros/OS levels	SLES 12 SP5 only	RHEL 8.6, SLES 15P3	Latest
Cluster	TCP/IP		EFA (RDMA equivalent)
I/O Fencing	Disk leasing (60+ secs)		SCSI-3PR equivalent (2-3 secs)
Cluster Size	Boxed in cluster size - XS, S, M, L		Potentially add XL
Cluster Manager	TSA		Pacemaker only

- New deployment will use 11.5.8.0 from here onwards
- Only the latest version within the same major release will be kept.

Q&A

